

JPEG2000 - Implications for Defence

Robert Prandolini, Mark Grigg and
Winston Fletcher

DSTO-TN-0408

DISTRIBUTION STATEMENT A:
Approved for Public Release -
Distribution Unlimited

20020531 102

JPEG2000 - Implications for Defence

Robert Prandolini, Mark Grigg and Winston Fletcher

**Information Technology Division
Electronics and Surveillance Research Laboratory**

DSTO-TN-0408

ABSTRACT

JPEG2000 is a competitive state-of-the-art new international image coding standard. Imagery is used by the Australian Defence Organisation (ADO) in many business processes covering strategic, operational and tactical levels. With the increasing international trend for the military to depend on international standards where appropriate, it is likely that JPEG2000 will be used in many military imagery systems and processes. The purpose of this report is to present an introduction to JPEG2000 and its implications for the ADO. The report consists of a background section on the present options for image coding, a technical overview of the new JPEG2000 standard, details of the implications for ADO imagery applications, and a conclusion including possible future research activities. The intended readers for this report are ADO personnel who are responsible for imagery technology.

RELEASE LIMITATION

Approved for public release

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE & TECHNOLOGY ORGANISATION | **DSTO**

AQ F02-08-1594

Published by

*DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567
© Commonwealth of Australia 2002
AR-012-124
January 2002*

APPROVED FOR PUBLIC RELEASE

JPEG2000 - Implications for Defence

Executive Summary

The goal of this report is to assist Defence in finding its technological directions such that the value of Australian Defence Organisation (ADO) imagery, and the productivity and effectiveness of ADO imagery systems and processes are maximised.

JPEG2000 is a new state-of-the-art international image coding standard. It is a significant improvement over the present JPEG image coding standard, has support from industry, and is likely to be widely adopted in many applications. This report presents an introduction to JPEG2000 and its implications for the ADO.

Defence is an environment that covers a large degree of scales, in many different parameters. Thus Defence requires an image compression system that scales well. JPEG2000 has been designed with scalability in mind and is well suited to Defence applications. It has a state-of-the-art compression performance, in both lossless and lossy modes; inherent multiscale resolution because it uses wavelet decomposition technology; tunable error-resilience capabilities; and an open embedded code-stream providing parsing related features.

With its adoption by industry, it is expected that JPEG2000 will become a common standard for imagery. This has two positive consequences. Firstly, interoperability is more likely for an international standard over proprietary standards. Secondly, it is likely that there will be multiple vendors of JPEG2000 based systems and services, creating an active and (hopefully) competitive marketplace.

The report discusses possible areas for future research in 1) minimising network and workstation requirements using intelligent JPEG2000 image servers, 2) extensions to the JPEG2000 standard for error-resilience techniques suitable for ADO tactical networks, and 3) the suitability of JPEG2000 for SAR and multi-spectral imagery.

Authors

Dr Robert Prandolini

Information Technology Division

Dr Robert Prandolini graduated with honours in Bachelor of Engineering from the Queensland University of Technology (QUT) in 1986. He worked as a tutor and then a lecturer at the QUT for over 5 years. During this time he was also engaged in a number of consultancy projects, the most interesting being in electronic forensics. This led to PhD studies, which he graduated from in 1996. He joined the Australian Federal Police (AFP) in 1995 as an expert in forensic electronic recording, audio analysis and enhancement. In 1999 he left the AFP to work on image processing for image coding and management for the Defence Science and Technology Organisation, Department of Defence, Australia.

Dr Mark Grigg

Information Technology Division

Dr Mark Grigg has a B.Eng (Electronics) and a B. App. Sci. (Physics) from RMIT, a PhD in Physics from the University of Melbourne and a Grad. Dip. Elec. Eng. from the UNSW. He joined DSTO in 1994 as a Research Scientist in ITD after working as a lecturer and tutor at various universities in Melbourne. Prior to that he worked as a consultant doing electronics and software development. He currently works with a team researching and developing prototype solutions to the challenges associated with the management and dissemination of imagery information in the ADO.

Mr Winston Fletcher

Information Technology Division

Winston Fletcher is an undergraduate student currently studying a Bachelor of Software Engineering degree at Swinburne University of Technology, Melbourne. He has completed three years of a 4 year course and is currently undertaking an optional 12 month industry placement at the Defence Science and Technology Organisation, Department of Defence. He is currently developing JPEG2000 client-server applications in the Java Programming language.

Contents

1. INTRODUCTION	1
2. BACKGROUND ON IMAGE CODING TECHNOLOGY	1
2.1 JPEG Standard	2
2.2 FlashPix	2
2.3 Wavelet Based Coders	3
2.3.1 LizardTech MrSID and ERMapper ECW	3
2.4 National Imagery Transmission Format (NITF)	3
3. DESIRABLE IMAGE CODING FEATURES	3
3.1 Summary of Defence Relevant Features of JPEG2000	4
3.2 State-of-the-art Compression	5
3.3 Open and Embedded Code-streams	6
3.4 Scalability	8
3.5 Error Resilience	8
4. THE IMPLICATIONS FOR DEFENCE OF JPEG2000	9
4.1 Status of the Standard	9
4.2 JPEG2000 Image Server	9
4.3 Multi-resolution and Multi-components	10
4.4 Error-resilience	10
5. CONCLUSIONS	11
APPENDIX A: JPEG2000 TECHNICAL OVERVIEW	13
A.1. Technical Description	13
A.1.1 Wavelet transform	13
A.1.2 Multiple components	14
A.1.3 Image tiles and transform code-blocks	14
A.1.4 JPEG 2000 code-stream syntax	15
A.1.5 Error resilience	16
6. BIBLIOGRAPHY	17

1. Introduction

Imagery is used by the Australian Defence Organisation (ADO) in many business processes including intelligence at the strategic, operational and tactical levels. There is a trend, both within Australia as well as internationally, towards the use of civilian standards for military purposes. JPEG2000 is a new competitive state-of-the-art international image coding standard. It is likely in the future that JPEG2000 will be used in many military imagery systems and processes. This report assesses the Defence requirements for image compression, compares these requirements with the capabilities offered by the JPEG2000 image compression standard, and finds this standard to be very capable at meeting Defence needs. It presents an introduction to JPEG2000 and its implications for Defence. The report consists of a background section on the present options for image coding, Defence requirements from image compression, JPEG2000 capabilities and details of its implications for ADO imagery applications. A technical overview of the new JPEG2000 standard is included in the appendix. The intended readers for this report are ADO personnel who are responsible for imagery technology.

2. Background on Image Coding Technology

Before discussing JPEG2000, a short introduction to the underlying concepts of image coding is presented in this section. Image data coding is the process of transforming the original representation of an image, most commonly an array of pixel values, into an alternate representation where the alternate representation is more compact than the original, hence the common term image compression. Image compression methods can be broadly categorised into two main classes. The first are the lossless techniques whereby there is no loss of information during the compression process. These techniques rely on coding redundancies in the image data to produce a reduction in the data file size. There are a number of lossless coding techniques, the most common of which are based on run length coding, dictionary coding or statistical coding such as the Huffman and arithmetic coding techniques.

The main advantage offered by the lossless coding techniques is the fact that there is no loss of any information in the compression process. The decompressed image data is identical to the original image data. Of course there are limits in the compression performance for lossless coding - the compression ratios achievable with these techniques are only moderate and rarely exceed about 2:1 or 3:1.

Lossy techniques on the other hand involve some loss of information from the original image. Using an appropriate compression technique, the information lost is usually visually redundant and thus the resultant image is generally indiscernible from the original. These techniques are most commonly based upon transform methods whereby the image data in the spatial domain, which is usually highly correlated, is transformed into a representation where the data can be decorrelated. This is very important as it allows each coefficient representing the image data to be treated independently without loss of compression efficiency. By applying knowledge of the human visual system and the fact that the response of the human visual system is very dependent upon spatial frequency, the transformed data can

effectively be reduced. This reduction is achieved by quantising the transformed data using visually weighted quantisation values. The process of quantising leads to an approximate representation, hence the loss of information, which in turn requires less data for the final representation of the compressed image.

It should be noted that the amount of compression that can be achieved by any given method depends very much on the characteristics of the image being compressed. While it is generally accepted that lossy schemes provide higher degrees of compression, they are not always the most appropriate for all image types where high compression rates are required. For example, graphics and diagrams are much better compressed using lossless techniques where high compression rates are achievable due to the nature of the images.

In the remainder of this section, we present a brief overview of the current JPEG standard coding scheme, the Flashpix format and two of the more common proprietary wavelet based coders. These are the MrSid and the Earth Resource Mappings ECW formats.

2.1 JPEG Standard

The most common of the lossy techniques, known as JPEG¹ (Joint Photographic Experts Group) has been established as an international standard for still image compression. This technique is based upon the use of the DCT (Discrete Cosine Transform) and is capable of producing high compression ratios without significant image degradation. Compression ratios of approximately 10:1 are achievable with virtually imperceptible image degradation and up to 100:1 or more if some image quality degradation is acceptable. The JPEG standard specifies a number of modes of operation which include, in addition to the most commonly used sequential baseline mode, progressive and hierarchical modes. The progressive mode allows for the delivery of images over transmission media such as the internet where a progressive refinement of the image occurs as more data is received. The hierarchical mode is useful in applications where a multiresolution representation is required.

2.2 FlashPix

The Flashpix format stores the images at multiple independent resolutions forming a hierarchical representation. Each image in the hierarchy is then divided into tiles of 64 pixels square. Each tile may be either uncompressed or, more commonly, compressed using JPEG compression. The final Flashpix compressed file size is approximately 1.5 times the size of a non-flashpix compressed file due to the hierarchical representation. The main advantage of the Flashpix format is the efficiency in the delivery of image data at any requested resolution level from an image server. Only the data necessary for the current viewing size and resolution is downloaded when the user navigates the image. This makes this system useful for the delivery of images over the internet.

¹ The JPEG standard also specifies a lossless compression technique that is not generally in common use.

2.3 Wavelet Based Coders

Wavelet based image coders have become available over the last few years. Wavelets offer a number of advantages over the standard JPEG coder, particularly at low compression data rates. The principles upon which wavelet coders are based are similar to the JPEG coder in that the image data is first decomposed into a frequency space representation. The wavelet coders can be designed to exploit the inherent multiresolution representation of the data giving the ability to decode images at various resolutions from a single data representation.

2.3.1 LizardTech MrSID and ERMapper ECW

MrSID (Multi resolution Seamless Image Database) by LizardTech and the ERMapper ECW format are examples of commercially available wavelet based compression systems. The features of these systems are similar, as they both use hierarchical wavelet representation to store the compressed image data. This representation of the image is an advantage over Flashpix because it does not incur any extra storage overhead. In addition, both are fully progressive with the ability to rapidly navigate an image by downloading only the necessary data for the current view size and resolution. Both the LizardTech and the ERMapper systems consist of a number of components providing an image server, viewer and plug-in components for various application packages. These systems are based on proprietary compression and delivery systems. Though this limits their adoption to anything other than supported applications, the range of supported viewers and plug-ins is extensive.

2.4 National Imagery Transmission Format (NITF)

The NITF, while strictly speaking not an image coding format, is important within the context of defence imagery systems as it represents a standard for the storage and dissemination of imagery products. NITFS consists of two parts, one relating to the representation of the image and associated information, and the second to the transmission of imagery data. The first part is basically a container for imagery consisting of a number of sections. These are a header, images, subimages, symbols, labels, textual information and other information related to the image in a user defined data section. The current NITF 2.1 standard [1] allows the following image types to be included in the images section of the file: uncompressed, bi-level, VQ (Vector Quantisation), JPEG, both lossy and lossless, correlated multicomponent compression algorithms and user defined compression. Work is currently underway to allow JPEG 2000 images to be included as part of the standard. These images will be included in the same way as the current images where JPEG 2000 images will be identified with a unique image category code in the NITF image header.

3. Desirable Image Coding Features

Within Defence, and in fact almost all environments that increasingly depend on ready access to many sources of information, there is a move to store imagery data exclusively in a digital form. This offers a number of advantages in terms of the storage, management and dissemination of this data that, paradoxically, are also its

disadvantages. Defence produces and consumes many types of imagery for various tasks – surveillance, reconnaissance, environmental monitoring, estate management, planning and logistics. For an example, consider surveillance imagery. It is produced with varying spatial resolution and size from a variety of platforms, and it can be arbitrarily large. Synthetic aperture radar (SAR) imaging systems can produce about 1 million pixels per second. Another example is the proliferation of handheld digital cameras, producing significant numbers of images suitable for intelligence, tactical and operational planning activities. These are but two examples of the large amount (in size and numbers) of imagery data finding applications in the ADO. Use of imagery is increasing at an exponential rate leading to storage and management problems and significant demands on network resources for dissemination. This is also creating the impetus for the procurement of image library and image management capabilities. Integral to storage and dissemination tasks is the image coding technology. Table 1 presents important image coding features that are relevant to Defence applications.

The image coding technology should have a state-of-the-art compression performance, be scalable in all permutations of the various parameters, and produce open embedded code-streams. That is, the technology should be able to support the extraction of the pertinent data from the code-stream of an image of arbitrary size, so that the decompressor can render an image; this image being an arbitrary segment of the whole image, at an arbitrary resolution, and with an arbitrary level of quality or precision. Further, it is possible to decode an embedded code-stream to any length to produce a resultant image with progressive quality - either progressive in precision, resolution, a mixture of both, or some other measure of quality. Such an embedded scalable image code-stream can support zooming and panning of image segments. Other desirable features include an optimal rate-distortion function, support for lossy and lossless compression, region of interest (ROI) transmission and fidelity priority, error resilience and interoperability. The following will present some discussion of these requirements and the performance of JPEG2000.

Table 1. Image coding requirements.

State-of-the-art compression
Embedded code-stream
Open architecture
Random access code-stream
Scalable - Image size
Scalable - Resolution
Scalable - Precision
Scalable - Image segment size
Optimal rate-distortion
Lossy and lossless compression
Region of interest
Error resilience

3.1 Summary of Defence Relevant Features of JPEG2000

Defence is an environment that covers a large degree of scales, in many different parameters. Thus Defence requires an image compression system that scales well.

JPEG2000 has been designed with scalability in mind and is well suited to Defence applications. It has a state-of-the-art compression performance, in both lossless and lossy modes; inherent multiscale resolution because it uses wavelet decomposition technology (so called reduced data sets are intrinsic in a typical JPEG2000 code-stream); tunable error-resilience capabilities; and an open embedded code-stream providing parsing related features (such as transcoding and client virtual image services). With its adoption by industry, it is expected that JPEG2000 will become a common standard for imagery. This has two positive consequences. Firstly, interoperability is more likely for an international standard over proprietary standards. Secondly, there is likely to be multiple vendors of JPEG2000 based systems and services, creating an active and (hopefully) competitive market place.

3.2 State-of-the-art Compression

Compression performance is of primary importance for Defence applications because the increasing number of images ranging up to very large sizes creates volumes that must be stored and transported. There are two types of modes for performance evaluation of compression technologies. The first is for high bit-rate lossless compression, or almost lossless compression. The second is for low bit-rate lossy compression.

True lossless compression means that the decompressed image is identical to the original image. This is called reversible compression. Lossless compression is used for archival purposes where it is important to store the image with no distortion. Lossless compression is also used when the image is to be processed through a number of stages that require decoding and re-encoding. This is because with lossy compression the accumulation of even minor distortions at each compression-decompression stage can result in significant coding artefacts.² While some compression methods are reversible, others are irreversible in that the decompressed image is not identical to the original. That said, such an image might only vary by a small amount so that it is almost lossless. Such image compression can be said to be perceptually lossless.

Comparisons between image compression technologies need to be evaluated at the lossless task and/or perceptually lossless task. This is considered high bit-rate compression, typically about a few bits per pixel (or pixel component).³ Its performance criterion is the compression rate achieved. The second evaluation mode is for low bit-rate compression (typically less than a bit per pixel) where the decompressed image may or may not have perceptual distortions. The criterion for comparison here is the image quality achieved for a given bit-rate.

At the moment, wavelet based image compression systems present the most competitive compression performance. JPEG2000 is a wavelet based compression system and is state-of-the-art. One of the objectives of the JPEG2000 project was improved performance for low bit-rate compression. JPEG2000 achieves a 20% to 50% improvement in compression performance over the JPEG standard. Comparisons can be found in [2-5]; and comparative demonstrations between

² This is so even though lossy compression can be used such that it is very close to lossless.

³ Typically 2 to 5 bits per pixel component, although this is dependent on the nature of the image.

JPEG2000 and JPEG can be found at [4, 6]. These comparisons show that JPEG2000 is significantly superior to JPEG for low bit-rate compression, as well as being superior to JPEG for high bit-rate compression.

Defence applications for expert image analysts and archival services would use lossless JPEG2000. Analyst product for intelligence publications and imagery services in operational and tactical environments would find lossy JPEG2000 more suitable. It should be re-emphasised that lossy coding can produce perceptually lossless imagery. The question of when to use lossless coding is based on whether the image is to be processed (analysed or enhanced) as some subsequent time after encoding.

3.3 Open and Embedded Code-streams

To encode an image at different rates using the JPEG standard, the compressor performs multiple coding passes to produce an independent code-stream per target rate. This requirement to re-encode per target bit-rate was one of the undesirable features of JPEG that the JPEG2000 project was established to address.

Embedded code-streams are defined as code-streams that can be decoded at any arbitrary length to produce an image. Thus the code-stream inherently contains the coding of the image at all bit-rates from 0 to some maximum value. This coding is also termed a one-pass compression method and its implications are threefold. The first is that in a single pass the encoder can code an image up to and including lossless or perceptually lossless quality. The second is that an image server or its client is able to truncate the code-stream at any length and still decode an image. Obviously, the image resolution and quality improves as more data is received, which leads to the third implication of progressive image display. The image can be decoded and displayed progressively from the code-stream as it is received, rather than having to wait for the entire code-stream to be received. This is an advantage for Defence applications that use low bandwidth networks.

Progressive transmission that allows images to be reconstructed with increasing pixel accuracy or spatial resolution is essential for many applications. Decoding and decompressing from an embedded code-stream allows the progressive reconstruction of the image with either progressively improving image resolution, or pixel accuracy, or some other quality schema. The delivery of an image to a tactical user with low bandwidth connectivity is one military example where a progressive quality display would be useful. Delivering images for the World Wide Web is another more common example.

It is desirable to have an open architecture code-stream so that different image types and applications can be optimised for their implementation. A possible Defence example for the future might be a transcoding service that acts as a bridge between terrestrial and wireless communication networks. This is more readily realisable with an open architecture code-stream. Also, an open architecture in standards assists in promoting multiple commercial vendors with interoperable (non-proprietary) services and interfaces. Thus open architectures can assist Defence acquisition projects to obtain better value for money.

One way of compartmentalising tasks for an image coding system is into compressed image code-stream generator (compressor), code-stream parser, and image decompressor. It is desirable if the compressor writes an image code-stream once (single pass compression), rather than having to re-write different code-streams for different compression parameters. A parser is able to read the code-stream and extract data as required. This data is sent to a decompressor that can produce an image from this data. Another task that the parser can perform is to recode the data into a new code-stream - a simple example of transcoding. A benefit of an open code-stream is that the parser task is relatively simple technology, which can then support scalable implementations of the decompressor. For example, a decompressor for an application that has a limited display size need only decode the specific data from the code-stream relevant to the task of rendering an image for that size display. The parser pulls out of the code-stream relevant data, based on viewing segment and resolution parameters, and then passes that data onto the decompressor. Now the point here is that such an implementation does not require the decompressor to be able to process images any larger than its target display. With the aid of the parser, it is able to render an appropriate representation of any image, even if that image was larger than the target display size. Thus the implementation for the parser and decompressor tasks is clean and simple. This is possible because the code-stream architecture is open and capable of supporting this implementation design.

Open embedded code-streams support random access into the data by the parser. This feature in turn supports user defined region of interest (ROI) capability. ROIs can be compressed with higher quality than the rest of the image and encoded earlier into the code-stream. Random code-stream processing can also assist more complex transcoding operations such as cropping, rotation, translation, filtering, feature extraction and scaling.

JPEG2000 implements the Embedded Block Coding with Optimised Truncation (EBCOT) [2] algorithm. It has an embedded and open architecture code-stream with a data structure that permits the above desirable features. A feature of JPEG2000 is that image manipulation such as cropping, rotation and translation can be performed in the compressed domain. JPEG2000 employs a layering of compressed image data in bins as shown in Figure 1. These data bins are placeholders for code-block bit-stream data. Code-block bit-streams are embedded and consist of the compressed image data relevant to a particular segment and resolution from the image. In Figure 1, the first quality layer (layer 1) has some data in the shown code-blocks' bins (numbered 1 to 7), except for code-block 3 which is empty. Note that the amount of data (represented by the height of the bars) is different for each code-block bin. The JPEG2000 code-stream consists of the concatenation of the quality layers to form a quality progressive embedded code-stream. All that a parser need do to support a decompressor is to read the data bins to produce the code-block bit-streams. The parser then sends the relevant code-block bit-streams to the decompressor for processing. One can clearly see how this code-stream design is open and embedded. Further, it supports scalability and random access into the code-stream data. More technical details of JPEG2000 are discussed in the appendix.




Figure 1. Progressive quality layers.
After Taubman [2].

3.4 Scalability

Scalability can be interpreted a number of different ways. It may mean that the technology is applicable (i.e. effective) at various levels of scales: small to large images, low to high image resolutions, low to high image quality, low to high compression bit-rates or low to high bandwidths. It may also mean that the implementation can be scalable, that is a compressor or decompressor may instance a subset of the technology capabilities. JPEG2000 provides scalability for both interpretations. The compression technology allows for a large range of scale in image size, resolutions, and compression bit-rates etc. The open architecture code-stream also permits scalable implementations. Cut-down or limited decompressors are defined by profiles in the standard, where the intention is to define performance requirements for standard profile decompressors that are targeted at applications with defined target limits (image size for example). Defence imagery types form a large range of sizes (i.e. height/width and resolution). Defence network types have a large range of bandwidth performance. So clearly, Defence requires a scalable image coding system.

3.5 Error Resilience

Communications in the tactical domain would have to be one of the more difficult environments, exhibiting poor channel performance - wireless communications frequently suffer from poor bit-error-rates. It is therefore desirable to consider the effect of bit-errors on the compressed image code-stream. Code-stream data dependencies are an important consideration. That is, portions of the code-stream may be more important than others in determining decoded image quality when the code-stream suffers bit-errors. The design of the code-stream can aid subsequent error correction systems in alleviating catastrophic decoding failures by having the data compartmented into independent bins of data. Data bins that have wide implications to the decompressor's performance should have more resilience to errors. Data bins that produce small localised distortions when in error obviously have a lower requirement for error resilience. Thus it is advantageous for an image coding system if it has appropriate error resilience. The JPEG2000 project considered error-resilience for the code-stream. Details of the error-resilience and a comparison of its superior performance over JPEG can be found in [3].

4. The Implications for Defence of JPEG2000

4.1 Status of the Standard

The JPEG2000 standard is being developed by the IOC/IEC Joint Technical Committee 1, Subcommittee 29, Working Group 1. This Working Group published the JPEG2000 Part I Final Draft International Standard [7] in September 2000, which then became the International Standard in December 2000. This release covers the compression technology and the code-stream syntax. While Part I is mandatory for compliance, Part II addresses features which are optional extensions for coding processes, code-stream syntax and file formats. Part II is scheduled for completion in November 2002. Overviews of JPEG2000 can be found in [3, 8, 9].

There are several other parts still to be standardised over 2001 and 2002. Part IV details conformance. There is agreement in the standards working group to create a number of technology profiles in Part IV. Their objective is to assist hardware manufacturers to design and produce JPEG2000 devices by removing some of the technology options found in Part I. Unfortunately, the suggested profiles [10] were either non-implementable, faulty in conception or removed important features from the Part I technology. The proposals would also remove the scalability for a number of parameters, thus creating non-interoperable versions of the JPEG2000 standard. This posed the significant risk that the lowest technology profile could become a de facto standard. Such an outcome would have harmed the availability of COTS JPEG2000 components suitable for Defence imagery. It was clearly in Defence's interests to assist in promoting a workable alternative for the draft Part IV profiles.

DSTO supported a move by the Australian National Body [11] to address the problems inherent with the Part IV draft. The outcome was the successful rewriting of the Part IV draft into a proposal that better addresses interoperability. The Australian effort has greatly reduced the risk of creating a de facto standard, while allowing some compromises in requirements for hardware manufactures. DSTO continues to be active in monitoring and participating in the development of the other parts of the JPEG2000 standard.

The standards committee has had significant and active participation from industry. Contributors include HP, Adobe, Sony, Ricoh, Sharp, Ericsson, Xerox, Eastman Kodak, Motorola, IBM, Microsoft and Canon. The original requirements of JPEG2000 [12] addressed the limitations of the present JPEG standard. A fair assessment is that the standard has achieved these requirements and will find acceptance in the community because of its superior technical qualities and strong industry support. At the time of writing, commercial JPEG2000 software was available from [13-16].

4.2 JPEG2000 Image Server

The JPEG2000 standard does not enunciate any methods for serving JPEG2000 images. Images can be served as JPEG2000 files, which are denoted by the file extension "jp2". In many domestic applications, such as web browsing, "jp2" files will be served in the same manner as the current JPEG "jpg" files.

However, Defence applications that use large images may not perform adequately with "jp2" image files in deployed operational or tactical environments. Limited available bandwidth, poor network performance and limited capacity of the client computer can eliminate the option of serving an entire large image file. Even in strategic environments, such as a headquarters, network bandwidth and computer workstation performance can place limits on the efficient use of imagery in applications. One solution to these restrictions would be to use client-side virtual images. A virtual image is where the client can view an image or a section of the image, at any desired resolution, as if the image data was stored locally. Since the client only displays a relatively small set of the data from the original image at any one time, a JPEG2000 server would serve the required image data as a JPEG2000 file on demand. It is known that Microsoft is presently developing a technology demonstrator of this type of JPEG2000 server [17].

Future JPEG2000 servers are likely to exploit client caching strategies so that data can be reused from previous views into the virtual image. Such intelligent servers would manage the data transmission based on a server-client session. The benefit of this method would be the minimisation of the transmitted data, which is an important consideration in Defence applications. DSTO has built a demonstration system of an intelligent JPEG2000 image server. It is expected that similar JPEG2000 image servers would be available commercially some time in the future. The Kakadu developers' software is available from [13] and includes JEG2000 web server and client software.

4.3 Multi-resolution and Multi-components

An important implication for Defence applications is that image coding algorithms that use wavelet transforms have inherent multi-resolution representations. Presently, an image analyst produces a reduced data set as part of the process for working with large images. This process may take some minutes to complete and requires an expensive workstation to process and store the data. It is a step that would not be required if JPEG2000 was used, as the image data is (typically) already in a Mallat decomposition [18]. Furthermore, an intelligent JPEG2000 server would be able to serve the data using smaller optimal data transactions. This would reduce the requirements on local network bandwidth and memory capacity for analyst workstations. JPEG2000 may bring about changes to the technical architecture of image analyst's workstations.

While JPEG2000 supports up to 16384 components, implementing JPEG2000 for multi-spectral and SAR imagery is an area of potential research. Part I describes the default transforms while Part II has extensions for explicit transform definitions. These extensions may provide a technical means to coding these special types of imagery. Consideration should be given for future investigations into image compression using JPEG2000 for multi-spectral and SAR imagery.

4.4 Error-resilience

An area of weakness for the JPEG2000 code-stream is the lack of error-resilience for header data. If an error occurs in this part of the packet, the result can be deleterious to the decoding of the entire image. The standard does not address this problem.

Strong error-resilience is required for Defence applications with high error-rate communication networks. Studies on the use of UDP and TCP for image transmission have shown how moderate channel bit-error-rates can significantly affect image data throughput [19]. What is required to address this problem is forward error correction (FEC). FEC adds redundancy to the bit-stream. It increases the amount of data and consequently reduces compression performance. But since the header information is small compared to the amount of coded coefficient data, header data FEC would only cause slight degradation in image compression performance. Thus header data FEC is highly desirable and has minimal overhead costs. The significance of the compressed coefficient data in a packet is dependent on a number of parameters, such as the importance of a region (ROI), the frequency sub-band and the bit-plane the data is from. Thus it would also be desirable to apply FEC to more important data, while less important data could rely on the "error-resilient segmentation marker" to detect an error.

The judicious use of FEC can be used to optimise image transmission performance under poor network conditions without significant cost to image quality performance. The standard includes "informational markers" for comments and it is therefore feasible to insert FEC into the code-stream as comments. A Defence application JPEG2000 encoder and decoder could be acquired with this FEC feature. Such a system would produce compliant JPEG2000 code-streams. A standard JPEG2000 decoder would not be able to use the FEC coded in the comments, but it would still be able to decode this JPEG2000 compliant code-stream (without the benefit of the FEC of course). Alternatively, a modified decoder could include a FEC processor that would correct packet header and significant-data errors. It is feasible that this could be done in a web browser by adding a FEC-comment parser plug-in antecedent to a standard JPEG2000 decoder plug-in. Thus the difficulty and costs for implementing FEC-comments are considered to be minimal. An area of future research to be considered is the dynamic optimisation of the application of FEC under Defence network environments.

5. Conclusions

The goal of this report is to assist Defence in finding its technological directions such that the value of ADO imagery, and the productivity and effectiveness of ADO imagery systems and processes is maximised. JPEG2000 is a competitive state-of-the-art image-coding standard. It will become the dominant international standard with support from industry and the users community. With the increasing international trend for the military to depend on international standards where appropriate, it is likely that JPEG2000 will be used in many military imagery systems and processes. Its future inclusion as part of the NITFS confirms its importance to Defence imagery systems. The required features for image compression have been presented in this report and they have been compared with those of JPEG2000. JPEG2000 has been found to be suitable for Defence applications. Its benefits and implications for Defence have been discussed.

Section 4.2 has outlined the possibility for minimising network and workstation requirements using intelligent JPEG2000 image servers. It is expected that such servers will become commercially available. Though scalability is implicit in the

JPEG2000 standard, it would be important to the ADO to help promote scalable capabilities with vendors. Defining performance scopes for the management of imagery in the ADO is the next step towards an image library and dissemination capability acquisition project.

Coding methods for SAR and multi-spectral imagery is an area of active research. This includes investigations into multiple-component transforms, wavelet packets, and other wavelet and non-wavelet based coding methods. JPEG2000 Part II has provisions for customisable transform kernel, customisable decomposition style, and multi-component transformations. It extends the flexibility of JPEG2000 Part I. It is proposed in Section 4.3 that future research on SAR and multi-spectral imagery compression should be considered. This research would not be limited to JPEG2000, however, what could be investigated is whether the most competitive compression methods available are adaptable for the JPEG2000 framework. Section 4.3 also foreshadowed possible implications from JPEG2000 on the technical architecture of image analyst's IT environment.

Section 4.4 has discussed an extension that uses FEC error-correcting code. This extension is outside of the JPEG2000 standard, but it is envisaged that such an error-resilient coder can be easily built using JPEG2000. Also raised is the use of UDP as the transport mechanism rather than TCP, and appropriate caching strategies for tactical applications. Research is proposed in JPEG2000 image servers for tactical applications. This would be important for supporting imagery applications in the tactical domain, which presently is considered to be unfeasible for Australian tactical communications systems.

Appendix A: JPEG2000 Technical Overview

A.1. Technical Description

JPEG2000 is an international standard being developed by the IOC/IEC Joint Technical Committee 1, Subcommittee 29, Working Group 1. Part I of the standard was published as an international standard in December 2000. It details the mandatory requirements for compliance to the standard. Part II of the standard, which is due for release in 2002, defines optional extensions. This section will give a brief overview of technical aspect of the JPEG2000 Part I standard. More information can be obtained from the tutorial papers [3, 8, 9] and presentation [4] available from Christopoulos's web page. Comparative performance results between JPEG2000 and JPEG can be found in [4] and [5]; and comparative demonstrations can be found at [4] and [6].

JPEG2000 is a wavelet based image coder. The algorithm used for compression is the Embedded Block Coding with Optimal Truncation (EBCOT) [2] by David Taubman. EBCOT features resolution and SNR scalability, independent embedded coding of blocks of coefficient data from the wavelet sub-bands, flexible packet wavelet decomposition, feature-rich random-accessible bit-stream, flexible prioritisation of coefficient bit-plane coding in significance layers, and competitive state-of-the-art performance.

A.1.1 Wavelet transform

The JPEG2000 encoder applies a Discrete Wavelet Transform (DWT) to the image. The most common wavelet decomposition is the Mallat [18] structure as shown in Figure A1. Here the original image is high-pass and low-pass filtered in the vertical and horizontal directions, and then down-sampled by 2. The three high-pass filtered sets of coefficients are denoted LH (vertical high-pass), HL (horizontal high-pass), and HH (vertical and horizontal high-pass). The low-pass filtered set of coefficients (LL) is the original image at a dyadic resolution. It may be filtered again to produce another set of LH, HL, HH and LL coefficients. In the figure, the filtering and downsampling produced 4 sets of coefficients at dyadic resolutions. The lowest resolution is level 0 (LL_0), followed by level 1 (LH_1, HL_1, HH_1), level 2 (LH_2, HL_2, HH_2), and level 3 (LH_3, HL_3, HH_3). Figure A2 shown the "Spac1" decomposition, which is another common wavelet packet decomposition.

Part I defines the three types of wavelets used to provide reversible and irreversible transforms for lossless and lossy coding. Part II [20] includes extensions that permit users to define their own transforms. The $w5 \times 3$ transform allows for lossless encoding (reversible) - it will produce an identical image to the original when decoded. Lossless coding can be used for archival purposes and for images that will be later processed - such that they will be decoded and re-encoded a number of times. The $w9 \times 7$ transform produces lossy coding. While lossy coding can be perceptually lossless, significant arithmetic rounding errors can accumulate if the image is decoded and re-encoded a number of times.




Figure A1. Mallat Decomposition. (After Taubman [2])




Figure A2. "Spac1" Wavelet Packet Decomposition. (After Taubman [2])

A.1.2 Multiple components

JPEG2000 files can encode images with 1 to a maximum of 16384 components and with a precision depth of up to 38 bits per pixel component. Each component has a wavelet transform applied to it independently. The standard permits different transforms to be applied to different components.

For colour images, the standard includes a reversible and irreversible transform for use with RGB and YCbCr components. The purpose of these transforms is to improve compression for colour images. The standard also allows for chrominance component sub-sampling.

A.1.3 Image tiles and transform code-blocks

The standard permits tiling of the image. The default is to use 1 tile that is the size of the image. A DWT is independently applied to the components in each tile. This produces wavelet coefficients grouped into sub-bands for each tile. The default is to

decompose the image into the Mallat structure, however, the standard permits any wavelet packet decomposition. These sub-bands are partitioned into rectangular code-blocks as shown in Figure A3.

The width and height of the code-blocks are limited to powers of two with a minimum size of 4, a maximum size of 256, the default is 64 and the product of the width and height must not exceed 4096. The size of the code-blocks is the same for all sub-bands at all resolution levels. The coefficients from each code-block are entropy encoded by an arithmetic coder. This is done in 3 fractional bit-plane coding passes per bit-plane. The data from the coding passes is arranged into layers, where the layers are groupings of these bit-plane coding passes. Layers can have any number of coding passes, ranging from zero to the entire coding of the block. Each layer monotonically improves the image fidelity for the resolution level the coefficient data is from. The number of layers present is signalled in a header at the beginning of the file and all tiles have the same number of layers at each resolution level.

This arrangement allows for decoding images at different resolution levels in a resolution progressive manner. Furthermore, as the code-blocks are encoded independently, regions of interest (ROI) can be encoded at higher bit-rates and coded in earlier layers such that the ROI appears with higher quality and earlier in a progressive display.

A.1.4 JPEG 2000 code-stream syntax

JPEG 2000 files are comprised of a main header, a tile header for each tile in the image and the compressed data for the tile. The main header, which is approximately 100 bytes on average, contains details describing the dimensions of the original image, the number of components, the number of tiles and tile size, coding style attributes (such as the size of code-blocks, wavelet transform used and number of layers), and details of ROI. The main header always appears at the beginning of a JPEG 2000 file.

The data for each tile consists of a tile header and the compressed data for the tile. The tile header is present to identify the tile. It also provides the option of overriding parameter values set by the main header (for example a tile may have different coding style attributes).

The compressed data for the tile appears after the tile header. This data appears in




Figure A3. Sub-band code-blocks.

'packets'. Each packet consists of a packet header followed by the packet data. Packets represent information from a single component, image resolution level and code-block layer and may contain data from the code-blocks for all the sub-bands in a resolution level. The order that the packets appear in the bit stream is indicated in the main or tile header and all packets must be present even if there is no data present in the packet. (Note that an empty packet can be indicated by a single byte.)

Average packet header size ranges from 5 to 12 bytes long and include the following information:

- Zero length packet bit.
- Code block inclusion.
- The number of missing most significant bit-planes in the code-blocks.
- The number of coding passes for each code-block in a packet.
- Length of the code-block data

The headers are stored in a compressed way using a scheme called tag trees, which is an efficient way of representing a two dimensional array of numbers. A bit stuffing routine is also used such that packets are byte-aligned by stuffing the last byte with zeros. Furthermore if the value of a byte is 0xFF, the next byte is stuffed with an extra 0 bit. This is to remove the possibility for ambiguity with the optional two-byte markers (0xFF91 and 0xFF92) that may be placed around the packets headers. The packet bodies are always comprised of a whole number of bytes and contain the data for included code-blocks.

A.1.5 Error resilience

A JPEG2000 code-stream consists of header information and compressed coefficient data, which is coded into packets. An error within a packet (the compressed data) will produce some distortion in the decoded image, however, the effects of this error is limited to the code-block that was corrupted. This is because blocks are coded independently. Thus the effects do not propagate outside of this block and the net distortion is limited to the frequency sub-band and spatial location of the affected code-block. To further limit the consequences of the distortion, the standard includes an "error-resilient segmentation marker" that will detect the occurrence of an error in the compressed coefficient data. Further error-resilient features include the termination of the arithmetic coder and the resetting of its context after each coding pass. This can be used to limit the distortion from an error to the one coding pass that the data error is from in the code-stream.

The following three allowable modes are particularly useful for error-resilience. The standard permits bypassing the arithmetic coder to encode the raw coefficient data. The standard is flexible and allows packet headers to be moved to special file header packets. The raw data can then be placed in short packets with resynchronisation markers. These features can be used to implement standard error-resilience techniques such as those found in [21-23]. Note that it has been found that short packets are important for error-resilience [19]. To utilise these error-resilient features included in the standard, the JPEG2000 encoder would have to incorporate a scheme for the implementation of error-resilience based on network performance. It is likely that this would not be available in a "domestic" encoder, however, it is expected that coders from specialist commercial vendors would be available. This is because one

of the significant application drivers behind the strength of error-resilience in JPEG2000 is wireless internet communication. A standard compliant decoder would be able to decode the code-stream and make use of these error-resilient features.

6. Bibliography

- [1] NIMA, "National Imagery Transmission Format V2.2 MIL-STD-2500B Notice 2,"., 2001.
- [2] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Transactions on Image Processing*, vol. 9, no. No. 7 July, pp. 1158-1170, 2000.
- [3] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 Still Image Coding System: An Overview," *IEEE Transactions on Consumer Electronics*, vol. 46, no. 4, pp. 1103-27, November 2000.
- [4] C. A. Christopoulos, A. N. Skodras, and T. Ebrahimi, "JPEG2000: The Next Generation Still Image Coding System,". Porto, Portugal, 2000, pp. 359-366.
- [5] M. Charrier, D. Santa Cruz, and M. Larsson, "JPEG2000, the Next Millennium Compression Standard for Still Images," presented at ICMCS, Florence, Italy, 7-11 June, 1999.
- [6] D. Santa Cruz, "Visual Demos,"
<http://ltswww.epfl.ch/~neximage/decoder/applets/>, 1999
- [7] M. Boliek, C. Christopoulos, and E. Majani, "JPEG 2000 Part I Final Draft International Standard," ISO/IEC JTC 1/SC 29/WG 1 N1890R, 25 September 2000.
- [8] M. Rabbani and R. Joshi, "An Overview of the JPEG2000 Still Image Compression Standard," *Signal Processing: Image Communication*, vol. 17, no. 1, Oct 2001.
- [9] A. Skodras, C. Christopoulos, and T. Ebrahimi, "The JPEG 2000 still image compression standard," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 36-58, 2001.
- [10] B. Brower, "JPEG 2000 Part 4 CD Release,".: ISO/IEC JTC1/CS29/WG1, 2000.
- [11] Australian National Body, "Australian National Body Comment on JPEG2000 Part 4 CD (WG1N2010),".: ISO/IEC JTC1/SC29/WG1, 2001.
- [12] Requirements ad hoc group, "Requirements and features in JPEG2000 Part I," ISO/IEC JTC 1/SC 29/WG 1 N1434, September 1999.
- [13] D. Taubman, "Kakadu: A software framework for JPEG2000 developers,"
<http://maestro.ee.unsw.edu.au/~taubman/kakadu/>, 2001
- [14] Algo_Vision, <http://www.algovision.com>,
- [15] Aware, <http://www.aware.com>,
- [16] Lead_Technologies, <http://www.leadtools.com>,
- [17] J. Li, "Personal communications with Prandolini,"., 2000.
- [18] S. G. Mallat, "Multifrequency channel decompositions of images and wavelet models," *IEEE Transactions on Acoustics, Speech, & Signal Processing*, vol. 37, no. 12, pp. 2091-110, 1989.
- [19] R. Prandolini, T. A. Au, A. K. Lui, M. J. Owen, and M. W. Grigg, "Use of UDP for efficient imagery dissemination," presented at Visual Communications and Image Processing, Perth, Australia, 20-23 June, 2000.
- [20] ISO/IEC JTC1/SC29/WG1, M. Boliek, C. Christopoulos, and E. Majani, "JPEG 2000 Part II Working Draft Version 1.0," ISO/IEC JTC 1/SC 29/WG 1 N1649, 14 April 2000.

- [21] M. W. Grigg, "Error Resilient Coding Technique Applied to Transmission of Imagery over HF Radio," presented at International Picture Coding Symposium, Melbourne, 1996.
- [22] J. Liang and R. Talluri, "Tools for robust image and video coding in JPEG2000 and MPEG4 standards," presented at Visual Communications and Image Processing '99. . SPIE. Soc. Imaging Sci. & Technol. 25-27 Jan. 1999., San Jose, CA, 25-27 Jan, 1999.
- [23] I. Moccagatta, S. Soudagar, J. Liang, and H. Chen, "Error-resilient coding in JPEG-2000 and MPEG-4," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 899-914, 2000.

JPEG2000 - Implications for Defence

Robert Prandolini, Mark Grigg and Winston Fletcher

(DSTO-TN-0408)

AUSTRALIA

DEFENCE ORGANISATION

Task Sponsor

DISREW

S&T Program

Chief Defence Scientist)	
FAS Science Policy)	
AS Science Corporate Management)	1 shared copy
Director General Science Policy Development)	
Counsellor Defence Science, London (Doc Data Sheet)		
Counsellor Defence Science, Washington (Doc Data Sheet)		
Scientific Adviser to MRDC Thailand (Doc Data Sheet)		
Scientific Adviser Joint		
Navy Scientific Adviser (Doc Data Sheet and distribution list only)		
Scientific Adviser - Army (Doc Data Sheet and distribution list only)		
Air Force Scientific Adviser		
Director Trials		
Director Aeronautical and Maritime Research Laboratory		
Director Electronics and Surveillance Research Laboratory (Doc Data Sheet and distribution list only)		
Chief Information Technology Division (Doc Data Sheet)		
Research Leader Command & Control and Intelligence Systems (Doc Data Sheet)		
Research Leader Military Information Enterprise		
Head, C3 Information Systems Concepts Group		
Head, Enterprise Visualisation, Instrumentation and Synchronisation Group (Doc Data Sheet)		
Head, Systems Simulation and Assessment Group (Doc Data Sheet)		
Head, C3I Operational Analysis Group (Doc Data Sheet)		
Head, Information Exploitation Group (Doc Data Sheet)		
Head, Intelligence Group (Doc Data Sheet)		
Head, Human Systems Integration Group (Doc Data Sheet)		
Head, C2 Australian Theatre		
Head, Distributed Systems Group (Doc Data Sheet)		
Dr Mark Grigg (Task Manager) C3ISC Fernhill		
Authors:		
Robert Prandolini		
Mark Grigg		
Winston Fletcher		
Publications & Publicity Officer, ITD /Executive Officer, ITD		1 shared copy

DSTO Library and Archives

Library Salisbury, 2 copies
US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy
National Library of Australia, 1 copy

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet)
Director General Land Development (Doc Data Sheet)
Director General Aerospace Development (Doc Data Sheet)

Knowledge Staff

Director General Command, Control, Communications and Computers
(DGC4) (Doc Data Sheet only)

Navy

SOSCI Surface Combatants FEG Wartime Division, Maritime HQ,
BLD 95, Garden Island, Locked Bag 12, Pyrmont NSW 2009 (Doc Data
Sheet)

Army

ABCA National Standardisation Officer, Land Warfare Development Sector
Tobruk Barracks, Puckapunyal VIC 3662, 4 copies
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), MILPO
Gallipoli Barracks, Enoggera QLD 4052 (Doc Data Sheet only)
NPOC QWG Engineer NBCD Combat Development Wing, Tobruk Barracks,
Puckapunyal, 3662 (Doc Data Sheet relating to NBCD matters only)

Intelligence Program

DGSTA Defence Intelligence Organisation
Manager, Information Centre, Defence Intelligence Organisation
Director, Defence Imagery and Geospatial Organisation

Corporate Support Program

Library Manager, DLS-Canberra

Universities and Colleges

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong, 3217
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

Other Organisations

NASA (Canberra)
AusInfo
State Library of South Australia

OUTSIDE AUSTRALIA

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US
Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and Technology, US

Spares (5 copies)

Total number of copies: 50

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE JPEG2000 - Implications for Defence			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Robert Prandolini, Mark Grigg and Winston Fletcher			5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Edinburgh SA 5111 Australia		
6a. DSTO NUMBER DSTO-TN-0408	6b. AR NUMBER AR-012-124		6c. TYPE OF REPORT Technical Note	7. DOCUMENT DATE January 2002	
8. FILE NUMBER N9505/23/10	9. TASK NUMBER JNT99/004	10. TASK SPONSOR DISREW	11. NO. OF PAGES 22	12. NO. OF REFERENCES 17	
13. URL ON WORLD WIDE WEB http://www.dsto.defence.gov.au/corporate/reports/DSTO-TN-0408.pdf			14. RELEASE AUTHORITY Chief, Information Technology Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CASUAL ANNOUNCEMENT Yes					
18. DEFTEST DESCRIPTORS Images Image processing Digital image processing					
19. ABSTRACT JPEG2000 is a competitive state-of-the-art new international image coding standard. Imagery is used by the Australian Defence Organisation (ADO) in many business processes covering strategic, operational and tactical levels. With the increasing international trend for the military to depend on international standards where appropriate, it is likely that JPEG2000 will be used in many military imagery systems and processes. The purpose of this report is to present an introduction to JPEG2000 and its implications for the ADO. The report consists of a background section on the present options for image coding, a technical overview of the new JPEG2000 standard, details of the implications to ADO imagery applications, and a conclusion including possible future research activities. The intended readers for this report are ADO personnel who are responsible for imagery technology.					

